JMLR: Workshop and Conference Proceedings 60:1–10, 2016 ACML 2016

Reinforcement-based Simultaneous Algorithm and its Hyperparameters Selection

Valeria Efimova EFIMOVA@RAIN.IFMO.RU
Andrey Filchenkov AFILCHENKOV@CORP.IFMO.RU
Anatoly Shalyto SHALYTO@MAIL.IFMO.RU
ITMO University, St. Petersburg, Kronverksky Pr. 49

arXiv:1611.02053v1 [cs.LG] 7 Nov 2016

Abstract

Many algorithms for data analysis exist, especially for classification problems. To solve a data analysis problem, a proper algorithm should be chosen, and also its hyperparameters should be selected. In this paper, we present a new method for the simultaneous selection of an algorithm and its hyperparameters. In order to do so, we reduced this problem to the multi-armed bandit problem. We consider an algorithm as an arm and algorithm hyperparameters search during a fixed time as the corresponding arm play. We also suggest a problem-specific reward function. We performed the experiments on 10 real datasets and compare the suggested method with the existing one implemented in Auto-WEKA. The results show that our method is significantly better in most of the cases and never worse than the Auto-WEKA.

Keywords: algorithm selection, hyperparameter optimization, multi-armed bandit, reinforcement learning

1. Introduction

The goal of supervised learning is to find a data model for a given dataset that allows to make the most accurate predictions. To build such model, lots of *learning algorithms* exist, especially in classification. These algorithms show various performances on different tasks. It prevents usage of a single universal algorithm to build a data model for all existing datasets. The performance of most of these algorithms depends on *hyperparameters*, the selection of which dramatically affects the performance of the algorithms.

Automated simultaneous selection of a learning algorithm and its hyperparameters is a sophisticated problem. Usually, this problem is divided into two subproblems that are solved independently: algorithm selection and hyperparameter optimization. The first is to select an algorithm from a set of algorithms (algorithm portfolio). The second is to find the best hyperparameters for preselected algorithm.

The first subproblem is typically solved by testing each of the algorithms with prechosen hyperparameters in the portfolio by many practitioners. Other methods are also in use, such as selecting algorithms randomly, by heuristics or using *k-fold cross-validation* (Rodriguez et al., 2010). But the last method requires running and then comparing all the algorithms. The other methods are not universally applicable. However, this subproblem has been in the scope of research interest for decades. Decision rules were used in several decades old papers

© 2016 V. Efimova, A. Filchenkov & A. Shalyto.

on algorithm selection from a portfolio (Aha, 1992). As an example, such rules are created to choose from 8 algorithms in (Ali and Smith, 2006).

Nowadays, more effective approaches exists such as meta learning (Giraud-Carrier et al., 2004; Abdulrahman et al., 2015). This approach is to reduce the algorithm selection problem to a supervised learning problem. It requires a training set of datasets D. For all $d \in D$, meta-feature vector is evaluated. Meta-features are useful characteristics of datasets, such as number of categorical or numerical features of an object $x \in d$, size of d and many others (Filchenkov and Pendryak, 2015; Castiello et al., 2005). After that, all the algorithms are run on all the datasets $d \in D$. Thus class labels are formed based on empirical risk evaluation. Then a meta-classifier is learnt on the prepared data with datasets as objects and best algorithms as labels. It is worth to note that it is better to solve this problem as the learning to rank problem (Brazdil et al., 2003; Sun and Pfahringer, 2013).

The second subproblem is a hyperparameter optimization that is to find hyperparameter vector for a learning algorithm that leads to the best performance of this algorithm for a given dataset. For example, hyperparameters of the Support Vector Machine (SVM) include kernel function and its hyperparameters; for a neural net, they include the number of hidden layers and the number of neurons in each of them. In practice, algorithms hyperparameters are usually chosen manually (Hutter et al., 2015). Moreover, sometimes the selection problem can be reduced to a simple optimization problem (primarily for statistical and regression algorithms), as, for instance, in (Strijov and Weber, 2010). However, this method is not universally applicable. Since hyperparameter optimization of classification algorithms is often applied manually, it requires a lot of time and do not lead to acceptable performance. There are several algorithms to solve the second subproblem automatically: Grid Search (Bergstra and Bengio, 2012), Random Search (Hastie et al., 2005), Stochastic Gradient Descent (Bottou, 1998), Tree-structured Parzen estimator (Bergstra et al., 2011), and the Bayesian Optimization including Sequential Model-Based Optimization (SMBO) (Snoek et al., 2012). In (Hutter et al., 2011), Sequential model-based algorithm configuration (SMAC) is introduced. It is based on SMBO algorithm. Another idea is implemented in predicting the best hyperparameter vector with meta-learning approach (Mantovani et al., 2015). Reinforcement-based approach was used in (Jamieson and Talwalkar, 2015) to operate several optimization threads with different settings.

Solution for the simultaneous selection of an algorithm and its hyperparameters is important for machine learning applications, but only a few of papers are devoted to this search. Moreover, these papers consider only a special case.

One of the possible solutions is to build a huge set of algorithms with prechosen hyperparameters and select from it. This solution was implemented in (Leite et al., 2012), in which a set of about 300 algorithms with chosen hyperparameters was used. However, such pure algorithm selection approach cannot provide any insurance of these algorithms quality for a new problem. This set may simply not include a hyperparameter vector for one of the presented learning algorithms with the best performance.

Another possible solution is sequential optimization of hyperparameters for every learning algorithm in portfolio and selection the best of them. This solution is implemented in the Auto-WEKA library (Thornton et al., 2012), it allows to choose one of the 27 base learning algorithms, 10 meta-algorithms and 2 ensemble algorithms and optimize its hyperparameters with SMAC method simultaneously and automatically. This method is described in

detail in (Thornton et al., 2012). It is clear that if we use the method, then it takes enormous time and may be referred to as exhaustive search (while, in fact, it is not due to the infinity of hyperparameter spaces).

The goal of this work is to suggest a method for simultaneous learning algorithm and its parameters selection being faster than the exhaustive search without affecting found solution quality. In order to do so, we use multi-armed bandit-based approach.

The remainder of this paper is organized as follows. In Section 2, we describe in details the learning algorithm and its hyperparameter selection problem and its two subproblems. The suggested method, based on multi-armed bandit problem, is presented in Section 3. In Section 4, experiment results are presented and discussed. Section 5 concludes.

This paper extends a paper accepted to International Conference on Intelligent Data Processing: Theory and Applications 2016.

2. Problem Statement

Let Λ be a hyperparameter space related to a learning algorithm A. We will denote the algorithm with prechosen hyperparameter vector $\lambda \in \Lambda$ as A_λ.

Here is the formal description of the algorithm selection problem. We are given a set of algorithms with chosen hyperparameters $\mathcal{A} = \{A^1_{\lambda_1}, \dots A^m_{\lambda_m}\}$ and learning dataset $D = \{d_1, \dots d_n\}$, where $d_i = (x_i, y_i)$ is a pair consisting of an object and its label. We should choose a parametrized algorithm $A^*_{\lambda^*}$ that is the most effective with respect to a quality measure Q. Algorithm efficiency is appraised by the use of dataset partition into learning and test sets with the further empirical risk estimation on the test set.

$$Q(A_\lambda, x) = \frac{1}{|D|} \sum_{x \in D} L(A_\lambda, x),$$

where $L(A_\lambda, x)$ is a loss function on object x, which is usually $L(A_\lambda, x) = [A_\lambda(x) \neq y(x)]$ for classification problems.

The algorithm selection problem thus is stated as the empirical risk minimization problem:

$$A^*_{\lambda^*} \in \operatorname*{argmin}_{A^j_{\lambda_j} \in \mathcal{A}} Q(A^j_{\lambda_j}, D).$$

Hyperparameter optimization is the process of selecting hyperparameters $\lambda^* \in \Lambda$ of a learning algorithm A to optimize its performance. Therefore, we can write:

$$\lambda^* \in \operatorname*{argmin}_{\lambda \in \Lambda} Q(A_\lambda, D).$$

In this paper, we consider the simultaneous algorithm selection and hyperparameters optimization. We are given learning algorithm set $\mathscr{A} = \{A^1, \dots, A^k\}$. Each learning algorithm A^i is associated with hyperparameter space Λ^i. The goal is to find algorithm $A^*_{\lambda^*}$ minimizing the empirical risk:

$$A^*_{\lambda^*} \in \operatorname*{argmin}_{A^j \in \mathscr{A}, \lambda \in \Lambda^j} Q(A^j_\lambda, D).$$

We assume that hyperparameter optimization is performed during the sequential hyperparameter optimization process. Let us give formal description. *Sequential hyperparameter optimization process* for a learning algorithm A^i:

$$\pi_i(t, A^i, \{\lambda_j^i\}_{j=0}^k) \to \lambda_{k+1}^i \in \Lambda^i.$$

It is a hyperparameter optimization method run on the learning algorithm A^i with time budget t, also it stores best found hyperparameter vectors within previous k iterations $\{\lambda_j\}_{j=0}^k$.

All of the hyperparameter optimization methods listed in the introduction can be described as a sequential hyperparameter optimization process, for instance, Grid Search or any of SMBO algorithm family including SMAC method, which is used in this paper.

Suppose that a sequential hyperparameter optimization process π_i is associated with each learning algorithm A_i. Then the previous problem can be solved by running all these processes. However, a new problem arises, the best algorithm search time minimization problem. In practice, there is a similar problem that is more interesting in practical terms. It is the problem of finding the best algorithm by fixed time. Let us describe it formally.

Let T be a time budget for the best algorithm $A_{\lambda^*}^*$ searching. We should split T into intervals $T = t_1 + \cdots + t_m$ such that if we run process π_i with time budget t_i we will get minimal empirical risk.

$$\min_j Q(A_{\lambda_j}^j, D) \xrightarrow[(t_1,\ldots,t_m)]{} \min,$$

where $A^j \in \mathscr{A}, \lambda_j = \pi_j(t_j, A^j, \emptyset)$ and $t_1 + \ldots + t_m = T$; $t_i \geq 0 \, \forall i$.

3. Suggested method

In this problem, the key source is a hyperparameter optimization time limit T. Let us split it up to q equal small intervals t and call them *time budgets*. Now we can solve time budgets assignment problem. Lets have a look at our problem in the different way. For each time interval, we should choose a process to be run during this interval before this interval starts.

The quality that will be reached by an algorithm on a given dataset is a priori unknown. On the one hand, the time spent for searching hyperparameters of not the best learning algorithms is subtracted from the time spent to improve hyperparameters for the best learning algorithm. On the other hand, if the time will be spent only for tuning single algorithm, we may miss better algorithms. Thus, since there is no marginal solution, the problem seems to be to find a tradeoff between exploration (assigning time for tuning hyperparameters of different algorithms) and exploitation (assigning time for tuning hyperparameters of the current best algorithm). This tradeoff detection is the classical problem in reinforcement learning, a special case of which is multi-armed bandit problem (Sutton and Barto, 1998). We cannot assume that there is a hidden process for state transformation that affects performance of algorithms, thus we may assume that the environment is static.

Multi-armed bandit problem is a problem, in which there are N bandit's arms. Playing each of the arms grants a certain reward. This reward is chosen according to an unknown probability distribution, specific to this arm. At each iteration k, an agent chooses an arm a_i and get a reward $r(i, k)$. The agent's goal is to minimize the total loss by time T. In this paper, we use the following algorithms solving this problem (Sutton and Barto, 1998):

1. ε-greedy: on each iteration, average reward $\bar{r}_{a,t}$ is estimated for each arm a. Then the agent plays the arm with maximal average reward with probability $1-\varepsilon$, and a random arm with probability ε. If you play each arm an infinite number of times, then the average reward converges to the real reward with probability 1.

2. UCB1: initially, the agent plays each arm once. On iteration t, it plays arm a_t that:

$$a_t \in \underset{i=1..N}{\operatorname{argmax}} \overline{r_{i,t}} + \sqrt{\frac{2 \cdot \ln t}{n_i}},$$

   where $\overline{r_{i,t}}$ is an average reward for arm i, n_i is the number of times arm i was played.

3. Softmax: initially, the agent plays each arm once. On iteration t, it plays arm a_i with probability:

$$p_{a_i} = \frac{e^{\bar{r}_i/\tau}}{\sum_{j=1}^{N} e^{r_j/\tau}},$$

   where τ is positive temperature parameter.

In this paper, we associate arms with sequential hyperparameters optimization processes $\{\pi_i(t, A^i, \{\lambda_k\}_{k=0}^q) \to \lambda_{q+1}^i \in \Lambda^i\}_{i=0}^m$ for learning algorithms $\mathscr{A} = \{A^1, \ldots, A^m\}$. After playing arm $i = a_k$ at iteration k, we assign time budget t to a process π_{a_k} to optimize hyperparameters. When time budget runs out, we receive hyperparameter vector λ_k^i. Finally, when selected process stops, we evaluate the result using empirical risk estimate for process π_i at iteration k, that is $Q(A^i_{\lambda_k^i}, D)$.

The algorithm we name MASSAH (**M**ulti-**a**rmed **s**imultanous **s**election of **a**lgorithm and its **h**yperparameters) is presented listing 1. There, MABSolver is implementing a multi-armed bandit problem solution, getConfig(i) is a function that returns $A^i_{\lambda_q}$, which is the best found configuration by q iterations to algorithm A^i.

The question we need to answer is how to define a reward function. The first (and simplest) way is to define a reward as the difference between current empirical risk and optimal empirical risk found during previous iterations. However, we meet several disadvantages. When the optimization process finds hyperparameters that leads to almost optimal algorithm performance, the reward will be extremely small. Also, the selection of such a reward function does not seem to be a good option for MABs, since probability distribution will depend on the number of iterations.

In order to find a reward function, such that the corresponding probability distribution will not change during the algorithm performance, we apply a little trick. Instead of defining reward function itself, we will define an average reward function. In order to do so, we use SMAC algorithm features.

Let us describe SMAC algorithm. At each iteration, a set of current optimal hyperparameter vectors is known for each algorithm. A local search is applied to find hyperparameter vectors which have distinction in one position with an optimal vector and improve algorithm quality. These hyperparameter vectors are added to the set. Moreover, some random hyperparameter vectors are added to the set. Then selected configurations (the algorithms with their hyperparameters) are sorted by *expected improvement* (EI). Some of the best configurations are run after that.

Algorithm 1: MASSAH
Data: D is the given dataset q is the number of iterations,
t is time budget for one iteration,
$\{\pi_i\}_{i=1,\dots,N}$ are sequential hyperparameter optimization processes.
Result: A_λ is algorithm with chosen hyperparameters

for $i = 1, \dots, N$ **do**
  $\lambda_i \leftarrow \pi_i(t, A^i, \emptyset)$
  $e_i \leftarrow Q(A^i_{\lambda_i}, D)$
end
$best_err \leftarrow \min_{i=1,\dots,N} e_i$
$best_proc \leftarrow \operatorname{argmin}_{i=1,\dots,N} e_i$
for $j = 1, \dots, q$ **do**
  $i \leftarrow$ MABSolver$(\{\pi_i\}_{i=1,\dots,N})$
  $\lambda_i \leftarrow \pi_i(t, A^i, \{\lambda_k\}_{k=1}^j)$
  $e_i \leftarrow Q(A^i_{\lambda_i}, D)$
  if $e_i < best_err$ **then**
    $best_err \leftarrow e_i$
    $best_proc \leftarrow i$
  end
end
return getConfig(π_{best_proc})

As in SMAC, we use empirical risk expectation at iteration k: $E_t(Q(A^i_{\lambda^i_k}, D))$, where $Q(A^i_{\lambda^i_k}, D)$ is empirical risk value reached by process π_i on dataset D at iteration k.

Note that process π_i optimizes hyperparameters for empirical risk minimization, but a multi-armed bandit problem is maximization problem. Therefore, we define an average reward function as:

$$\bar{r}_{i,(k)} = \frac{Q_{max} - E_{(k)}(Q(A^i_{\lambda^i_k}, D))}{Q_{max}},$$

where Q_{max} is the maximal empirical risk that was achieved on a given dataset.

4. Experiments

Since Auto-WEKA implements the only existing solution, we choose it for comparison. Experiments were performed on 10 different real datasets with a predefined split into training and test data from UCI repository[1]. These datasets characteristics are presented in Table 1.

The suggested approach allows to use any hyperparameter optimization method. In order to perform comparison properly, we use SMAC method that is used by Auto-WEKA. We consider 6 well-known classification algorithms: k Nearest Neighbors (4 categorical and 1 numerical hyperparameters), Support Vector Machine (4 and 6), Logistic Regression (0 and 1), Random Forest (2 and 3), Perceptron (5 and 2), and C4.5 Decision Tree (6 and 2).

1. http://www.cs.ubc.ca/labs/beta/Projects/autoweka/datasets/

Table 1: Datasets description.

Dataset	Number of categorical features	Number of numerical features	Number of classes	Number of objects in training set	Number of objects in test set
Dexter	0	20000	2	420	180
German Credit	13	7	2	700	300
Dorothea	0	100000	2	805	345
Yeast	0	8	10	1039	445
Secom	0	590	2	1097	470
Semeion	0	256	10	1116	477
Car	6	0	4	1210	518
KR-vs-KP	36	0	2	2238	958
Waveform	0	40	3	3500	1500
Shuttle	38	192	2	35000	15000

As we previously stated, we are given time T to find the solution of the main problem. The suggested method requires splitting T into small equal intervals t. We give the small interval to a selected process π_i at each iteration. We compare the method performance for different time budget t values to find the optimal value. We consider time budgets from 10 to 60 seconds with 3 second step. After that we run the suggested method on 3 datasets Car, German Credits, KRvsKP described above. We use 4 solutions of the multi-armed bandit problem: UCB1, 0.4-greedy, 0.6-greedy, Softmax. We run each configuration 3 times. The results show no regularity, so we assume time budget t as 30 seconds.

In the quality comparison, we consider suggested method with the different multi-armed bandit problem solutions: UCB1, 0.4-greedy, 0.6-greedy, Softmax with the naïve reward function, and two solutions $UCB1_{E(Q)}$, $Softmax_{E(Q)}$ with the suggested reward function. Time budget on iteration is $t = 30$ seconds, the general time limitation is $T = 3$ hours $=$ 10800 seconds. We run each configuration 12 times with random seeds of SMAC algorithm. Auto-WEKA is also limited to 3 hours and selects one of the algorithms we specified above. The experiment results are shown in Table 2.

The results show that the suggested method is significantly better in most of the cases than Auto-WEKA of the all 10 datasets, because its variations reach the smallest empirical risk. There is no fundamental difference between the results of the suggested method variations. Nevertheless, algorithms $UCB1_{E(Q)}$ $Softmax_{E(Q)}$, which use the suggested reward function, achieved the smallest empirical risk in most cases.

The experiment results show that the suggested approach improves the existing solution of the simultaneous learning algorithm and its hyperparameters selection problem. Moreover, the suggested approach does not impose restrictions on a hyperparameter optimization process, so the search is performed on the entire hyperparameters space for each learning algorithm. It is significant that the suggested method allows to select a learning algorithm with hyperparameters, whose quality is not worse than Auto-WEKA outcome quality.

We claim that the suggested method is statistically not worse than Auto-WEKA. To prove this, we carried out Wilcoxon signed-rank test. In experiments, we use 10 datasets

Table 2: Comparison of Auto-WEKA and suggested methods for selecting classification algorithm and its hyperparameters for the given dataset. We performed 12 independent runs of each configuration and report the smallest empirical risk Q achieved by Auto-WEKA and the suggested method variations. We highlight with bold entries that are minimal for the given dataset.

Dataset	AutoWEKA	UCB1	0.4-greedy	0.6-greedy	Softmax	$UCB1_{E(Q)}$	$Softmax_{E(Q)}$
Car	0.3305	**0.1836**	**0.1836**	**0.1836**	**0.1836**	**0.1836**	**0.1836**
Yeast	34.13	**29.81**	**29.81**	33.65	**29.81**	**29.81**	**29.81**
KR-vs-KP	0.2976	**0.1488**	**0.1488**	**0.1488**	**0.1488**	**0.1488**	**0.1488**
Semeion	4.646	**1.786**	**1.786**	**1.786**	**1.786**	**1.786**	**1.786**
Shuttle	**0.00766**	0.0115	0.0115	**0.00766**	0.0115	**0.0076**	**0.0076**
Dexter	7.143	2.38	2.381	2.381	2.381	2.381	**0.16**
Waveform	11.28	**8.286**	**8.286**	**8.286**	**8.286**	**8.286**	**8.286**
Secom	4.545	**3.636**	4.545	4.545	**3.636**	**3.636**	**3.636**
Dorothea	6.676	4.938	4.958	4.938	4.938	4.32	**2.469**
German Credits	19.29	**14.29**	**14.29**	15.71	**14.29**	**14.29**	**14.29**

which leads to an appropriate number of pairs. Moreover, other Wilcoxon test assumptions are carried. Therefore, we have 6 test checks: comparison of Auto-WEKA and each variation of the suggested method. Since the number of samples is 10, we have meaningful results when untypical results sum $T < T_{0,01} = 5$. We consider a minimization problem, so we test only the best of 12 runs for each dataset. Finally, we have $T = 3$ for the ε-greedy algorithms and $T = 1$ for the others. This proves the statistical significance of the obtained results.

5. Conclusions

In this paper, we suggest and examine a new solution for the actual problem of an algorithm and its hyperparameters simultaneous selection. The proposed approach is based on a multi-armed bandit problem solution. We suggest a new reward function exploiting hyperparameter optimization method properties. The suggested function is better than the naïve function in applying a multi-armed bandit problem solutions to solve the main problem. The experiment result shows that the suggested method outperforms the existing method implemented in Auto-WEKA.

The suggested method can be improved by applying meta-learning in order to evaluate algorithm quality to preprocess a given dataset before running any algorithm. This evaluation can be used as a prior knowledge of an algorithm reward. Moreover, we can add a context vector to hyperparameters optimization process and use solutions of a contextual multi-armed bandit problem. We can select some datasets by meta-learning and then get the empirical risk estimate and use it as context.

Acknowledgments

Authors would like to thank Vadim Strijov and unknown reviewers for useful comments. The research was supported by the Government of the Russian Federation (grant 074-U01) and the Russian Foundation for Basic Research (project no. 16-37-60115).

References

Salisu Mamman Abdulrahman, Pavel Brazdil, Jan N van Rijn, and Joaquin Vanschoren. Algorithm selection via meta-learning and sample-based active testing. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases; International Workshop on Meta-Learning and Algorithm Selection*. University of Porto, 2015.

David W Aha. Generalizing from case studies: A case study. In *Proc. of the 9th International Conference on Machine Learning*, pages 1–10, 1992.

Shawkat Ali and Kate A Smith. On learning algorithm selection for classification. *Applied Soft Computing*, 6(2):119–138, 2006.

James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *The Journal of Machine Learning Research*, 13(1):281–305, 2012.

James S Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems*, pages 2546–2554, 2011.

Léon Bottou. Online learning and stochastic approximations. *On-line learning in neural networks*, 17(9):142, 1998.

Pavel B Brazdil, Carlos Soares, and Joaquim Pinto Da Costa. Ranking learning algorithms: Using ibl and meta-learning on accuracy and time results. *Machine Learning*, 50(3): 251–277, 2003.

Ciro Castiello, Giovanna Castellano, and Anna Maria Fanelli. Meta-data: Characterization of input features for meta-learning. In *International Conference on Modeling Decisions for Artificial Intelligence*, pages 457–468. Springer, 2005.

Andrey Filchenkov and Arseniy Pendryak. Datasets meta-feature description for recommending feature selection algorithm. In *Artificial Intelligence and Natural Language and Information Extraction, Social Media and Web Search FRUCT Conference (AINL-ISMW FRUCT), 2015*, pages 11–18. IEEE, 2015.

Christophe Giraud-Carrier, Ricardo Vilalta, and Pavel Brazdil. Introduction to the special issue on meta-learning. *Machine learning*, 54(3):187–193, 2004.

Trevor Hastie, Robert Tibshirani, Jerome Friedman, and James Franklin. The elements of statistical learning: data mining, inference and prediction. *The Mathematical Intelligencer*, 27(2):83–85, 2005.

Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *Learning and Intelligent Optimization*, pages 507–523. Springer, 2011.

Frank Hutter, Jörg Lücke, and Lars Schmidt-Thieme. Beyond manual tuning of hyperparameters. *KI-Künstliche Intelligenz*, 29(4):329–337, 2015.

Kevin Jamieson and Ameet Talwalkar. Non-stochastic best arm identification and hyperparameter optimization. *JMLR*, 41:240–248, 2015.

Rui Leite, Pavel Brazdil, and Joaquin Vanschoren. Selecting classification algorithms with active testing. In *Machine Learning and Data Mining in Pattern Recognition*, pages 117–131. Springer, 2012.

Rafael Gomes Mantovani, André LD Rossi, Joaquin Vanschoren, André Carlos Ponce de Leon Carvalho, et al. Meta-learning recommendation of default hyper-parameter values for svms in classifications tasks. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases; International Workshop on Meta-Learning and Algorithm Selection*. University of Porto, 2015.

Juan D Rodriguez, Aritz Perez, and Jose A Lozano. Sensitivity analysis of k-fold cross validation in prediction error estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(3):569–575, 2010.

Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959, 2012.

Vadim Strijov and Gerhard Wilhelm Weber. Nonlinear regression model generation using hyperparameter optimization. *Computers & Mathematics with Applications*, 60(4):981–988, 2010.

Quan Sun and Bernhard Pfahringer. Pairwise meta-rules for better meta-learning-based algorithm ranking. *Machine learning*, 93(1):141–161, 2013.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 1998.

Chris Thornton, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Auto-weka: Automated selection and hyper-parameter optimization of classification algorithms. *CoRR, abs/1208.3719*, 2012.

This figure "spiral.png" is available in "png" format from:

http://arxiv.org/ps/1611.02053v1